arXiv:1610.04702v1 [math.OC] 15 Oct 2016

Optimal Distributed Stochastic Mirror Descent for Strongly Convex Optimization ⋆

Deming Yuan [a], Yiguang Hong [b], Daniel W. C. Ho [c], Guoping Jiang [a],

[a] *College of Automation, Nanjing University of Posts and Telecommunications, Nanjing 210023, Jiangsu, P.R. China*

[b] *Key Laboratory of Systems and Control, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, P.R. China*

[c] *Department of Mathematics, City University of Hong Kong, Kowloon, Hong Kong*

Abstract

In this paper we consider convergence rate problems for stochastic strongly-convex optimization in the non-Euclidean sense with a constraint set over a time-varying multi-agent network. We propose two efficient non-Euclidean stochastic subgradient descent algorithms based on the Bregman divergence as distance-measuring function rather than the Euclidean distances that were employed by the standard distributed stochastic projected subgradient algorithms. For distributed optimization of non-smooth and strongly convex functions whose only stochastic subgradients are available, the first algorithm recovers the best previous known rate of $O(\ln(T)/T)$ (where T is the total number of iterations). The second algorithm is an epoch variant of the first algorithm that attains the optimal convergence rate of $O(1/T)$, matching that of the best previously known centralized stochastic subgradient algorithm. Finally, we report some simulation results to illustrated the effectiveness of the proposed algorithms.

Key words: Distributed stochastic optimization; Strong convexity; Non-Euclidean divergence; Mirror descent; Epoch gradient descent; Optimal convergence rate

1 Introduction

Recent years have witnessed a growing interest in developing distributed subgradient algorithms for solving convex constrained optimization problem, where the objective function is the sum of the local convex objective functions of nodes in a network (see, e.g., Nedić & Ozdaglar, 2009; Gharesifard & Cortés, 2014; Lou, Shi, Johansson, & Hong, 2014; Zhu & Martínez, 2012; Lin, Ren, & Song, 2016), due to their widespread applications including sensor networks (see, e.g., Nedić & Olshevsky, 2016; Shi, Ling, Wu, & Yin, 2015), distributed machine learning (see, e.g., Duchi, Agarwal, & Wainwright, 2012; Jakovetić, Xavier, & Moura, 2014), and smart grid (see, e.g., Yi, Hong, & Liu, 2015; Chang, Nedić, & Scaglione, 2014), to name a few. The algorithms for solving the problem need to be fully distributed, in the sense that each node can only operate on its local objective function and exchange information with its instant neighbors.

Strong convexity has been widely studied in convex optimization, because strongly convex cost functions can be easily found in a variety of engineering application domains like sensor networks and smart grids and strongly convex properties are actively used in regularization methods. Take the ridge regression problem as an example, where the objective function consists of the strongly convex Tikhonov regularization term for some performance improvement in optimization computation (see, e.g., Shalev-Shwartz, & Ben-David, 2014). In light of the increasing attention to distributed optimization, various distributed designs for optimizing strongly convex functions (in the Euclidean sense) have been proposed in the literature (see Nedić & Olshevsky, 2016 and Tsianos & Rabbat, 2012), due to its wide application in many practical fields and its potential to provide better guarantees of convergence performance.

⋆ This paper was not presented at any IFAC meeting. This work was supported by the National Natural Science Foundation of China under Grant 61304042. Corresponding author D. Yuan. Tel. +86 25 85866512. Fax +86 25 85866512.

Email addresses: dmyuan1012@gmail.com (Deming Yuan), yghong@iss.ac.cn (Yiguang Hong), madaniel@cityu.edu.hk (Daniel W. C. Ho), jianggp@njupt.edu.cn (Guoping Jiang).

Many algorithms have been developed over the past years to solve distributed convex optimization problems (see, e.g., Nedić, Ozdaglar, & Parrilo, 2010; Lu, Tang, Regier, & Bow, 2011; Liu, Qiu, & Xie, 2014; Yuan, Ho, & Xu, 2016; Yuan, Ho, & Hong, 2016; Liu & Wang, 2015; Kia, Cortés, & Martínez, 2015; Ram, Nedić, & Veeravalli, 2010). Such algorithms require only the first-order information of the objective functions and Euclidean projection onto the constraint set. This makes the algorithms attractive for large-scale optimization problems. However, the aforementioned algorithms are inherently Euclidean, in the sense that they rely on measuring distances based on Euclidean norms. This means that it is challenging or infeasible to generate efficient projections for certain objective functions and constraint sets, taking the Euclidean projection onto the unit simplex as an example. In this paper, we shall develop a class of distributed algorithms that are built on mirror descent (also known as non-Euclidean subgradient descent; see, e.g., Nemirovski, Juditsky, Lan, & Shapiro, 2009; Beck & Teboulle, 2003), which generalizes the projection step using the Bregman divergence. Bregman divergences are a general class of distance-measuring functions, which include the Euclidean distance and Kullback-Leibler (KL) divergence as special cases. The work Xi, Wu, and Khan (2014) presents a study of the distributed optimization algorithm that builds on mirror descent, for solving the non-strongly and deterministic variant of problem (1); however, only asymptotic convergence results are established for the proposed algorithm.

Convergence rate is an important issue in the distributed design. Although the aforementioned algorithms in the last paragraph can be applied to distributed optimization of strongly convex functions, it is desirable to develop algorithms by further exploiting the strongly convexity of the objective function, in order to provide better performance such as *faster* convergence rates. In fact, establishing the convergence rates of the algorithms for strongly convex functions has become an active line of distributed optimization research. In Nedić and Olshevsky (2016), the authors proposed a distributed stochastic subgradient-push algorithm for solving problem (1), under the assumption that the stochastic gradients of the objective functions are Lipschitz. In particular, the algorithm converges at an $O(\ln(T)/T)$ rate in the unconstrained case, which is (to the best of our knowledge) the previously best known rate in the literature. The work Rabbat (2015) utilized the similar idea as that of Xi, Wu, and Khan (2014)) to develop a distributed proximal subgradient algorithm (that can be viewed as a Euclidean variant of the one in Xi, Wu, and Khan (2014)) for solving the unconstrained composite stochastic optimization problems; they prove that the proposed algorithm converge at an $O(1/T)$ rate, under the smoothness assumptions on the objective functions. Notably, the authors in Tsianos and Rabbat (2012) proposed a class of distributed algorithms (in both batch and online setting) that converge at an $O(\ln(T)/T)$ rate in the constrained case, without making the smoothness assumptions on the objective functions. The algorithm development borrows the idea of epoch gradient descent algorithms from the machine learning community, which are capable of accelerating the convergence rates by better exploiting the strongly convexity of the objective functions and utilizing the step sizes (see, e.g., Hazan & Kale, 2014). However, no conclusive investigation has been carried out using this idea to design and analyze the mirror descent based distributed algorithms.

In this paper we focus on establishing the convergence rate of algorithms for the distributed strongly convex constrained optimization problem in the following form

$$\begin{array}{ll} \text{minimize} & F(\mathbf{w}) = \sum_{i=1}^{m} F_i(\mathbf{w}) \\ \text{subject to} & \mathbf{w} \in \mathcal{W} \end{array} \tag{1}$$

where each F_i (which is only known to node i) is strongly convex in the non-Euclidean sense and maybe non-smooth, and $\mathcal{W} \subseteq \mathbf{R}^d$ is a convex constraint set known to all the nodes in the network. Moreover, the nodes can only compute the noisy subgradients of their respective objective functions. To be specific, we assume that there exists a stochastic subgradient oracle, which for any point $\mathbf{w} \in \mathcal{W}$, returns a random estimate $\hat{\mathbf{g}}_i(\mathbf{w})$ of a subgradient $\mathbf{g}_i(\mathbf{w}) \in \partial F_i(\mathbf{w})$ so that $\mathbb{E}[\hat{\mathbf{g}}_i(\mathbf{w})] = \mathbf{g}_i(\mathbf{w})$, where $\partial F_i(\mathbf{w})$ denotes the subdifferential set of $F_i(\cdot)$ at $\mathbf{w}$. It is well-known that for (centralized) stochastic optimization of non-smooth and strongly convex functions, the optimal convergence rate is $O(1/T)$ (see, e.g., Hazan & Kale, 2014; Rakhlin, Shamir, & Sridharan, 2012). This fact, combined with the above observations, motivates us to consider the following questions: 1) Is it possible to develop a distributed stochastic mirror descent algorithm that recovers the best previously known rate $O(\ln(T)/T)$, for distributed optimization of non-smooth and strongly convex functions? and 2) For the same optimization problem, is it possible to devise a variant of the developed algorithm that attains the optimal $O(1/T)$ convergence rate?

In this paper, we give affirmative answers to the above questions. Specifically, the main contributions of this paper are highlighted as follows:

- We consider the construction of non-Euclidean algorithms for distributed stochastic optimization of strongly convex functions (in the non-Euclidean sense) whose only stochastic subgradients are available. The algorithms generalize the standard distributed stochastic projected subgradient algorithms to the non-Euclidean setting. Therefore, the proposed algorithms are more flexible, in the sense that they enable us to generate efficient updates to better reflect the geometry of the underlying optimization problem,

by carefully choosing the Bregman divergence. As an important consequence, we apply the proposed algorithms to the stochastic strongly-convex optimization over the unit simplex. Specifically, we derive a distributed stochastic entropic descent algorithm, which can be given explicitly by simple formulas, as opposed to the standard Euclidean algorithms that always involve solving a d-dimensional nonlinear equation at every step.

- We propose a distributed stochastic mirror descent (DSMD) algorithm to answer the first question. In particular, we show that for a total number of T iterations, the proposed algorithm achieves an $O\left(\ln(T)/T\right)$ rate of convergence, by exploiting the strongly convexity of the objective functions. This is in contrast to the work Xi, Wu, and Khan (2014), where only asymptotic convergence is established. In addition, this rate recovers the best previous known rate in Nedić and Olshevsky (2016) and Tsianos and Rabbat (2012). Moreover, in contrast to the algorithm in Nedić and Olshevsky (2016), our proposed DSMD algorithm is in constrained setting, which naturally arises in a number of applications where each node's estimate has to lie within some decision space (see, e.g., Nedić, Ozdaglar, & Parrilo, 2010; Duchi, Agarwal, & Wainwright, 2012).
- We propose an epoch variant of the DSMD algorithm, called Epoch-DSMD algorithm, to answer the second question. The Epoch-DSMD algorithm combines the strength of the epoch gradient descent algorithm that is widely used in the machine learning community and the DSMD algorithm. Specifically, the Epoch-DSMD algorithm divides the total number of T iterations into a sequence of exponentially increasing epochs, and runs the DSMD algorithm on each epoch with a constant step size that is reduced by a constant factor from epoch to epoch. The starting point of each epoch is the average of the previous one. In particular, we prove by induction that the resulting point returned by the last epoch attains the optimal $O(1/T)$ rate of convergence, which largely improve the $O\left(\ln(T)/T\right)$ rate obtained by Tsianos and Rabbat (2012) with the Euclidean norm. To the best of our knowledge this is the first distributed algorithm that achieves this converge rate for distributed stochastic strongly-convex constrained optimization, without assuming the smoothness of the objective functions.

The reminder of the paper is organized as follows. In Section 2, we give a description of the network model that underlies problem (1), and provide the assumptions made on the problem. In Section 3, we propose a distributed stochastic mirror descent algorithm and provide its convergence analysis results. In Section 4, we extend the proposed algorithm to the epoch setting and establish its convergence analysis results. In Section 5, we report some simulations about the proposed algorithms. Finally, we conclude with Section 6.

Notation: Let $\mathbf{R}^d$ be the d-dimensional vector space, and $\mathbf{R}_{++}$ be the set of positive real numbers. Write $\|\mathbf{w}\|_2$ to denote the Euclidean norm of a vector $\mathbf{w} \in \mathbf{R}^d$, and $\langle \mathbf{w}, \mathbf{v} \rangle$ to denote the standard inner product on $\mathbf{R}^d$, for any $\mathbf{w}, \mathbf{v} \in \mathbf{R}^d$. We denote by $[m]$ the set of integers $\{1, \ldots, m\}$. For a vector $\mathbf{w}$, we denote its ith component by $[\mathbf{w}]_i$. We denote the (i,j)th element of a matrix $\mathbf{P}$ by $[\mathbf{P}]_{ij}$. For a differentiable function f, Let $\nabla f(\mathbf{w})$ denote the gradient of $f(\cdot)$ at $\mathbf{w}$, and $\mathbb{E}[X]$ denote the expected value of a random variable X.

2 Problem Setting and Assumptions

In this paper, we are interested in solving convergence rate problems for (1) over a time-varying multi-agent network. Specifically, let $\mathcal{G}(t) = (\mathcal{V}, \mathcal{E}(t), \mathbf{P}(t))$ be a directed graph that represents the nodes' communication pattern at time t, where $\mathcal{V} = \{1, \ldots, m\}$ is the node set, $\mathcal{E}(t)$ is the set of activated links at time t, and $\mathbf{P}(t)$ is the communication matrix at time t. We make the following standard assumption on graph $\mathcal{G}(t)$ (see, e.g., Ram, Nedić, & Veeravalli, 2010; Yuan, Ho, & Xu, 2016).

Assumption 1 *The graph $\mathcal{G}(t) = (\mathcal{V}, \mathcal{E}(t), \mathbf{P}(t))$ satisfies ($t = 1, 2, \ldots$):*

(a) There exist a scalar $0 < \xi < 1$ such that $[\mathbf{P}(t)]_{ii} \geq \xi$ for all i and t, and $[\mathbf{P}(t)]_{ij} \geq \xi$ whenever $(j,i) \in \mathcal{E}(t)$;
(b) $\mathbf{P}(t)$ is doubly stochastic, i.e., $\sum_{i=1}^m [\mathbf{P}(t)]_{ij} = 1$ and $\sum_{j=1}^m [\mathbf{P}(t)]_{ij} = 1$ for all i and j;
(c) There exists some positive integer B such that the graph $\left(\mathcal{V}, \bigcup_{t=sB+1}^{(s+1)B} \mathcal{E}(t)\right)$ is strongly connected for every $s \geq 0$.

An important consequence of Assumption 1 is given in the following lemma, adopted from Nedić, Olshevsky, Ozdaglar, and Tsitsiklis (2008), that characterizes the convergence properties of the transition matrices as follows:

$$\mathbf{P}(t, \ell) = \mathbf{P}(t)\mathbf{P}(t-1)\cdots\mathbf{P}(\ell), \qquad \forall t \geq \ell \geq 1.$$

Lemma 1 *Under Assumption 1, we have that, for all i, j and all $t \geq \ell \geq 1$, $\left|[\mathbf{P}(t,\ell)]_{ij} - \frac{1}{m}\right| \leq \alpha\beta^{t-\ell+1}$, with $\alpha = \left(1 - \frac{\xi}{4m^2}\right)^{-2}$ and $\beta = \left(1 - \frac{\xi}{4m^2}\right)^{1/B}$.*

To study non-Euclidean optimization, we first give the definition of the Bregman divergence, which is used as the distance-measuring functions in mirror descent, in contrast to the Euclidean distances employed by the standard distributed stochastic subgradient algorithms.

Definition 1 *Given a strongly convex and differentiable distance-generating function $\Phi : \mathbf{R}^d \to \mathbf{R}$, the* Bregman divergence *induced by Φ is defined as follows:*

$$D_\Phi(\mathbf{w}||\mathbf{v}) := \Phi(\mathbf{w}) - \Phi(\mathbf{v}) - \langle \nabla\Phi(\mathbf{v}), \mathbf{w} - \mathbf{v} \rangle .$$

Remark 1 *Here are the two standard choices of the distance-generating function Φ and its associated Bregman divergence D_Φ (see, e.g., Nemirovski, Juditsky, Lan, & Shapiro, 2009; Beck & Teboulle, 2003).*

(1) $\Phi(\mathbf{w}) = \frac{1}{2}\|\mathbf{w}\|_2^2$, *and the associated Bregman divergence is given by* $D_\Phi(\mathbf{v}||\mathbf{w}) = \frac{1}{2}\|\mathbf{v}-\mathbf{w}\|_2^2$;

(2) $\Phi(\mathbf{w}) = \sum_{i=1}^d [\mathbf{w}]_i \ln[\mathbf{w}]_i$, *and the associated Bregman divergence, also known as the Kullback-Leibler divergence, is as follows:*

$$D_\Phi(\mathbf{w}||\mathbf{v}) = \sum_{i=1}^d [\mathbf{w}]_i \ln \frac{[\mathbf{w}]_i}{[\mathbf{v}]_i}.$$

It is time to provide a non-Euclidean strong convexity assumption on problem (1).

Assumption 2 *(a) Φ is σ_Φ-strongly convex with respect to the Euclidean norm, where (without loss of generality) $\sigma_\Phi \geq 1$, i.e., for any two points $\mathbf{w}$, $\mathbf{v} \in \mathbf{R}^d$,*

$$\Phi(\mathbf{w}) \geq \Phi(\mathbf{v}) + \langle \nabla\Phi(\mathbf{v}), \mathbf{w}-\mathbf{v}\rangle + \frac{\sigma_\Phi}{2}\|\mathbf{w}-\mathbf{v}\|_2^2;$$

(b) $F_i(\mathbf{w})$ is σ_F-strongly convex with respect to function Φ, i.e., for any two points $\mathbf{w}$, $\mathbf{v} \in \mathcal{W}$, we have for any $i = 1, \ldots, m$,

$$F_i(\mathbf{w}) \geq F_i(\mathbf{v}) + \langle \mathbf{g}_i(\mathbf{v}), \mathbf{w}-\mathbf{v}\rangle + \sigma_F D_\Phi(\mathbf{w}||\mathbf{v})$$

where $D_\Phi(\mathbf{w}||\mathbf{v})$ are convex in their second argument $\mathbf{w}$ for every fixed $\mathbf{v}$.

Remark 2 *Note that Assumptions 2(a) is standard in developing mirror descent algorithms for solving convex optimization problems (see, for example, Nemirovski, Juditsky, Lan, & Shapiro, 2009; Beck & Teboulle, 2003). Assumption 2(b) is commonly used in developing distributed mirror descent algorithms for strongly convex optimization (see, e.g., Xi, Wu, & Khan, 2014; Li, Chen, Dong, & Wu, 2016). In fact, under Assumption 2(b), one can easily show that each F_i is also strongly convex with respect to the Euclidean norm, by simply assuming the smoothness of function Φ.*

The following assumption is about stochastic subgradient, which has been widely used in the literature (see, e.g., Nedić & Olshevsky, 2016).

Assumption 3 *At any point $\mathbf{w} \in \mathcal{W}$, the stochastic subgradient of function F_i satisfies:*

$$\mathbb{E}[\|\hat{\mathbf{g}}_i(\mathbf{w})\|_2^2] \leq G^2, \qquad \forall i.$$

3 Distributed Stochastic Mirror Descent

In this section, we propose a distributed stochastic algorithm that is built on mirror descent. The details of the non-Euclidean algorithm are given in Algorithm 1. At each iteration we utilize the Bregman divergence defined in Section 2 (rather than the Euclidean distances employed by most standard approaches) as the distance-measuring function. Specifically, we will establish the explicit convergence rate of the proposed algorithm. This is in contrast to the work Xi, Wu, and Khan (2014), where only asymptotic convergence is established.

Algorithm 1 Distributed Stochastic Mirror Descent (DSMD)

Input: total number of iterations T and step size sequence $\{\eta_t\}_{t=1}^T$

Initialize: $\mathbf{w}_{i,1} \in \mathcal{W}$ for all $i \in [m]$

1: **for** $t = 1$ to T **do**

2: Query the stochastic subgradient oracle at $\mathbf{w}_{i,t}$ to get a random $\hat{\mathbf{g}}_{i,t} := \hat{\mathbf{g}}_i(\mathbf{w}_{i,t})$ such that $\mathbb{E}[\hat{\mathbf{g}}_{i,t}] = \mathbf{g}_i(\mathbf{w}_{i,t})$ is a subgradient of F_i at $\mathbf{w}_{i,t}$

3: Compute

$$\begin{aligned}\nabla\Phi(\mathbf{v}_{i,t}) &= \nabla\Phi(\mathbf{w}_{i,t}) - \eta_t\hat{\mathbf{g}}_{i,t}\\ \mathbf{u}_{i,t+1} &= \arg\min_{\mathbf{w}\in\mathcal{W}} D_\Phi(\mathbf{w}||\mathbf{v}_{i,t})\\ \mathbf{w}_{i,t+1} &= \sum_{j=1}^m [\mathbf{P}(t)]_{ij}\mathbf{u}_{j,t+1}\end{aligned}$$

4: **end for**

Before presenting the main convergence results, we provide two standard setups for the distributed stochastic mirror descent algorithms, such discussions are applied to the Epoch-DSMD algorithm (see Section 4) as well.

- *Euclidean setup:* $\Phi(\mathbf{w}) = \frac{1}{2}\|\mathbf{w}\|_2^2$, and the associated Bregman divergence is $D_\Phi(\mathbf{v}||\mathbf{w}) = \frac{1}{2}\|\mathbf{v}-\mathbf{w}\|_2^2$. In that case Step 3 in Algorithm 1 recovers the following standard distributed stochastic projected subgradient algorithm:

$$\mathbf{u}_{i,t+1} = \Pi_{\mathcal{W}}\left(\mathbf{w}_{i,t} - \eta_t\hat{\mathbf{g}}_{i,t}\right) \tag{2}$$

$$\mathbf{w}_{i,t+1} = \sum_{j=1}^N [\mathbf{P}(t)]_{ij}\mathbf{u}_{j,t+1} \tag{3}$$

  where $\Pi_{\mathcal{W}}(\cdot)$ denotes the standard Euclidean projection onto the convex set $\mathcal{W}$.
- *Simplex setup:* For this setup, suppose that our constraint set $\mathcal{W} = \Delta_d = \{\mathbf{w} \in \mathbf{R}^d : \sum_{i=1}^d [\mathbf{w}]_i = 1, [\mathbf{w}]_i \geq 0, i \in [d]\}$. Let $\Phi(\mathbf{w}) = \sum_{i=1}^d [\mathbf{w}]_i \ln[\mathbf{w}]_i$, and it is proved that $\Phi(\mathbf{w})$ is 1-strongly convex with respect to the ℓ_1 norm $\|\cdot\|_1$ over the constraint set $\mathcal{W}$ (see, e.g., Nemirovski, Juditsky, Lan, & Shapiro, 2009; Beck & Teboulle, 2003). Its associated Bregman divergence is given in Remark 1, and it is straightforward to verify that $D_\Phi(\mathbf{w}||\mathbf{v})$ is convex in their second argument $\mathbf{v}$ for every fixed $\mathbf{w}$, i.e., Assumption

2(b) holds. More importantly, in that case Step 3 in Algorithm 1 is equivalent to the following *distributed stochastic entropic descent algorithm*:

$$[\mathbf{u}_{i,t+1}]_j = \frac{[\mathbf{w}_{i,t}]_j \exp(-\eta_t [\hat{\mathbf{g}}_{i,t}]_j)}{\sum_{\ell=1}^d [\mathbf{w}_{i,t}]_\ell \exp(-\eta_t [\hat{\mathbf{g}}_{i,t}]_\ell)} \quad j \in [d] \qquad (4)$$

$$\mathbf{w}_{i,t+1} = \sum_{j=1}^N [\mathbf{P}(t)]_{ij} \mathbf{u}_{j,t+1}. \qquad (5)$$

Noting that the update (2) in the standard distributed projected subgradient algorithm cannot be solved explicitly (in fact, it involves computing the solution of d-dimensional nonlinear equation at each step, as pointed out by Beck and Teboulle (2003)), as opposed to the update (4).

Now we establish the main result of the DSMD algorithm, via showing the convergence of the average point defined as follows:

$$\widehat{\mathbf{w}}_{i,T} = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_{i,t} \qquad \forall i \in [m]$$

to the optimal point $\mathbf{w}^* = \arg\min_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w})$.

Theorem 1 *Let Assumptions 1, 2, and 3 hold. Let $\eta_t = \frac{1}{\sigma_F t}$ for all $t = 1, \ldots, T$. Then, for any $j \in [m]$ and $T \geq 3$, we have*

$$\mathbb{E}\left[\|\widehat{\mathbf{w}}_{j,T} - \mathbf{w}^*\|_2^2\right] \leq \frac{c \ln(T)}{T} + \frac{c'}{T}$$

where

$$c = \frac{2G^2}{\sigma_F^2 \sigma_\Phi} \left(1 + \frac{4\alpha\beta m}{(1-\beta)\sigma_\Phi}\right) \qquad \text{and}$$

$$c' = \frac{2G}{\sigma_F \sigma_\Phi} \left(\frac{2\alpha\beta}{1-\beta} + 1\right) \sum_{i=1}^m \mathbb{E}[\|\mathbf{w}_{i,1}\|_2].$$

The proof of Theorem 1 relies on the following two lemmas. The first lemma establishes a bound on the differences among the estimates of all the nodes in the network, and the second lemma states the basic convergence result of Algorithm 1.

Lemma 2 (Disagreement) *Let Assumptions 1, 2, and 3 hold. Then*

$$\sum_{t=1}^T \sum_{i=1}^m \mathbb{E}[\|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2] \leq m \left(\frac{2\alpha\beta}{1-\beta} + 1\right) \sum_{i=1}^m \mathbb{E}[\|\mathbf{w}_{i,1}\|_2] + \frac{2\alpha\beta m^2 G}{(1-\beta)\sigma_\Phi} \sum_{t=1}^T \eta_t.$$

Proof. See Appendix A. □

Lemma 3 (Basic Convergence Result) *Let Assumptions 1 and 2 hold. Then, for any $t \in [T]$, we have*

$$\sum_{i=1}^m \langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{w}^* \rangle \leq \frac{\eta_t}{2} \sum_{i=1}^m \|\hat{\mathbf{g}}_{i,t}\|_2^2 + \sum_{i=1}^m \frac{D_\Phi(\mathbf{w}^* || \mathbf{w}_{i,t}) - D_\Phi(\mathbf{w}^* || \mathbf{w}_{i,t+1})}{\eta_t}.$$

Proof. See Appendix B. □

Armed with Lemmas 2 and 3, we are ready to present the proof of Theorem 1.

Proof. [Proof of Theorem 1] Let $\mathbb{E}_{|t-1}[X]$ denote the expectation conditioned on all the randomness until round $t-1$. By our assumption on the stochastic subgradient oracle $\hat{\mathbf{g}}_{i,t}$, we have $\mathbb{E}_{|t-1}[\hat{\mathbf{g}}_{i,t}] = \mathbf{g}_i(\mathbf{w}_{i,t})$. This fact implies

$$\sum_{i=1}^m \mathbb{E}_{|t-1} \langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{w}^* \rangle = \sum_{i=1}^m \langle \mathbf{g}_i(\mathbf{w}_{i,t}), \mathbf{w}_{i,t} - \mathbf{w}^* \rangle . \qquad (6)$$

Taking the total expectation and using Assumption 2(b) on the strongly convexity of F_i with respect to function Φ, we have

$$\begin{aligned}
&\sum_{i=1}^m \mathbb{E}[\langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{w}^* \rangle] \\
&= \sum_{i=1}^m \mathbb{E}[\langle \mathbf{g}_i(\mathbf{w}_{i,t}), \mathbf{w}_{i,t} - \mathbf{w}^* \rangle] \\
&\geq \sum_{i=1}^m \mathbb{E}\left[F_i(\mathbf{w}_{i,t}) - F_i(\mathbf{w}^*) + \sigma_F D_\Phi(\mathbf{w}^* || \mathbf{w}_{i,t})\right] \\
&\geq \sum_{i=1}^m \mathbb{E}[F_i(\mathbf{w}_{j,t})] - F(\mathbf{w}^*) + \sigma_F \sum_{i=1}^m \mathbb{E}[D_\Phi(\mathbf{w}^* || \mathbf{w}_{i,t})] \\
&\quad - G \sum_{i=1}^m \mathbb{E}[\|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2] \qquad (7)
\end{aligned}$$

where the last inequality is based on following relation that utilizes the convexity of function F_i:

$$\begin{aligned}
F_i(\mathbf{w}_{i,t}) &\geq F_i(\mathbf{w}_{j,t}) + \langle \mathbf{g}_i(\mathbf{w}_{j,t}), \mathbf{w}_{i,t} - \mathbf{w}_{j,t} \rangle \\
&\geq F_i(\mathbf{w}_{j,t}) - \|\mathbf{g}_i(\mathbf{w}_{j,t})\|_2 \cdot \|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2 \\
&\geq F_i(\mathbf{w}_{j,t}) - G\|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2 \qquad (8)
\end{aligned}$$

because $\|\mathbf{g}_i(\mathbf{w}_{j,t})\|_2 = \|\mathbb{E}[\hat{\mathbf{g}}_i(\mathbf{w}_{j,t})]\|_2 \leq \mathbb{E}[\|\hat{\mathbf{g}}_i(\mathbf{w}_{j,t})\|_2] \leq \left(\mathbb{E}[\|\hat{\mathbf{g}}_i(\mathbf{w}_{j,t})\|_2^2]\right)^{1/2} \leq G$. Combining inequality (7) with

the one in Lemma 3, we get

$$\begin{aligned}
&\sum_{t=1}^{T}\mathbb{E}[F(\mathbf{w}_{j,t})] - F(\mathbf{w}^*) \\
&\quad\le \sum_{t=1}^{T}\sum_{i=1}^{m}\frac{\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t})] - \mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t+1})]}{\eta_t} \\
&\qquad -\sigma_F\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t})] + \sum_{t=1}^{T}\frac{\eta_t}{2}\sum_{i=1}^{m}\mathbb{E}[\|\hat{\mathbf{g}}_{i,t}\|_2^2] \\
&\qquad +G\sum_{t=1}^{T}\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,t}-\mathbf{w}_{j,t}\|_2] \\
&\quad := p_1 + p_2 + p_3 + p_4
\end{aligned} \tag{9}$$

where p_1, p_2, p_3, p_4 denote each of the right-hand side terms in (9). For terms p_1 and p_2, using the fact that $\eta_t = \frac{1}{\sigma_F t}$ and the non-negativity of the Bregman divergence, we obtain

$$\begin{aligned}
p_1 + p_2 &= \left(\frac{1}{\eta_1} - \sigma_F\right)\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1})] \\
&\quad + \sum_{t=2}^{T}\left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} - \sigma_F\right)\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t})] \\
&\quad - \frac{1}{\eta_T}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,T+1})] \\
&= -\frac{1}{\eta_T}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,T+1})] \\
&\le 0.
\end{aligned} \tag{10}$$

Term p_3 can be bounded as follows, by using Assumption 3:

$$\sum_{t=1}^{T}\frac{\eta_t}{2}\sum_{i=1}^{m}\mathbb{E}[\|\hat{\mathbf{g}}_{i,t}\|_2^2] \le \frac{mG^2}{2}\sum_{t=1}^{T}\eta_t. \tag{11}$$

In addition, p_4 can be bounded by Lemma 2. Hence, combing Lemma 2 with inequalities (9), (10), and (11), and using the following two inequalities:

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\eta_t = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{\sigma_F t} &\le \frac{1}{\sigma_F}\cdot\frac{1}{T}\left(1 + \int_1^T \frac{1}{u}\mathrm{d}u\right) \\
&\le \frac{2}{\sigma_F}\cdot\frac{\ln(T)}{T}, \qquad \forall T \ge 3
\end{aligned}$$

and

$$\begin{aligned}
&\frac{1}{T}\sum_{t=1}^{T}F\left(\mathbf{w}_{j,t}\right) - F(\mathbf{w}^*) \\
&\ge F\left(\frac{1}{T}\sum_{t=1}^{T}\mathbf{w}_{j,t}\right) - F(\mathbf{w}^*) \\
&\ge \langle \mathbf{g}(\mathbf{w}^*), \widehat{\mathbf{w}}_{j,T} - \mathbf{w}^*\rangle + m\sigma_F D_\Phi(\widehat{\mathbf{w}}_{j,T}||\mathbf{w}^*) \\
&\ge m\sigma_F D_\Phi(\widehat{\mathbf{w}}_{j,T}||\mathbf{w}^*) \\
&\ge \frac{m\sigma_F\sigma_\Phi}{2}\|\widehat{\mathbf{w}}_{j,T} - \mathbf{w}^*\|_2^2
\end{aligned}$$

where the third line follows from the first order optimality condition, therefore, we complete the proof. □

Remark 3 *Problem (1) has a natural connection with the distributed online convex optimization, and the DSMD algorithm can be applied to solving the online variant of problem (1) as well. To simplify the presentation we suppose that each node i is endowed with a sequence of cost functions $\{F_{i,1}, F_{i,2}, \cdots, F_{i,T}\}$ that are drawn i.i.d. from some unknown distribution, where $F_{i,t}$ ($t \in [T]$) is the cost function of node i at time t, and the goal of the online variant of problem (1) is to solve the following optimization problem:*

$$\min_{\mathbf{w}\in\mathcal{W}}\sum_{t=1}^{T}\sum_{i=1}^{m}F_{i,t}(\mathbf{w}).$$

By simply defining the expected cost function as the objective function in problem (1), that is, $F_i(\mathbf{w}) = \mathbb{E}[F_{i,t}(\mathbf{w})]$, it is straightforward to show that the regret of the online variant of DSMD algorithm is $O(\ln(T))$.

Remark 4 *Theorem 1 recovers the best previously known rate of $O(\ln(T)/T)$ in Tsianos and Rabbat (2012), for distributed optimization of non-smooth and strongly convex functions. It is worth noting that our algorithm generalizes the one in Tsianos and Rabbat (2012) to non-Euclidean setting and the underlying network of our algorithm is time-varying. In contrast to Nedić and Olshevsky (2016), we do not impose the smoothness assumptions on the objective functions, and our algorithm is in constrained setting.*

In fact, we can adopt the reduction technique as that of Yuan, Ho, and Hong (2016) to derive an $O(1/\sqrt{T})$ convergence rate for the non-strongly variant of problem (1). The basic idea is to add a controlled amount of strong convexity to each objective function, and then apply the DSMD algorithm to the new objective function.

4 The Epoch-based Optimal Algorithm

In the previous section, we proposed a distributed stochastic mirror descent algorithm to achieve a rate of convergence at $O(\ln(T)/T)$, which is suboptimal for stochastic strongly-convex optimization. In this section, we present an optimal distributed stochastic mirror descent algorithm, called Epoch-DSMD, to solve problem (1) and analyze its convergence properties.

The epoch idea to improve optimization algorithms was proposed in the study of machine learning (see Hazan & Kale, 2014). Based on its two-stage structure, we have more freedom of degrees to adjust algorithms for better performance. To our knowledge, the idea was not applied to the convergence rate analysis for distributed optimization problems, and the Epoch-DSMD is a first try to propose a distributed algorithm that combines the ideas of epoch gradient descent algorithms in the machine learning studies with those of the mirror descent in non-Euclidean optimization. The details of the Epoch-DSMD algorithm are shown in Algorithm 2. Similar to the epoch gradient descent algorithm in Hazan and Kale (2014), the Epoch-DSMD consists of a sequence of epochs, each of which has an exponentially increasing iteration number and an exponentially decreasing step size (Step 8). In each epoch, the updates are just the DSMD algorithm (Step 2 to Step 5), and each node only maintains the average point and use it as the starting point of the next epoch (Steps 6 and 7).

Algorithm 2 Epoch-DSMD

Input: an initial step size η_1, number of iterations in the first epoch T_1, and total number of iterations T

Initialize: $\mathbf{w}_{i,1}^1 = \arg\min_{\mathbf{w}\in\mathcal{W}} \Phi(\mathbf{w})$ for all $i \in [m]$, and set $k = 1$

1: **while** $\sum_{\ell=1}^{k} T_\ell \leq T$ **do**
2: **for** $t = 1$ to T_k **do**
3: Query the stochastic subgradient oracle at $\mathbf{w}_{i,t}^k$ to get a random $\hat{\mathbf{g}}_{i,t}^k := \hat{\mathbf{g}}_i(\mathbf{w}_{i,t}^k)$ such that $\mathbb{E}[\hat{\mathbf{g}}_{i,t}^k] = \mathbf{g}_i(\mathbf{w}_{i,t}^k)$ is a subgradient of F_i at $\mathbf{w}_{i,t}^k$
4: Compute

$$\nabla\Phi(\mathbf{v}_{i,t}^k) = \nabla\Phi(\mathbf{w}_{i,t}^k) - \eta_k \hat{\mathbf{g}}_{i,t}^k$$
$$\mathbf{u}_{i,t+1}^k = \arg\min_{\mathbf{w}\in\mathcal{W}} D_\Phi(\mathbf{w}||\mathbf{v}_{i,t}^k)$$
$$\mathbf{w}_{i,t+1}^k = \sum_{j=1}^{m} [\mathbf{P}_k(t)]_{ij} \mathbf{u}_{j,t+1}^k$$

5: **end for**
6: Compute $\widehat{\mathbf{w}}_i^k = \frac{1}{T_k}\sum_{t=1}^{T_k} \mathbf{w}_{i,t}^k$
7: Update $\mathbf{w}_{i,1}^{k+1} = \widehat{\mathbf{w}}_i^k$
8: Update $T_{k+1} = 2T_k$ and $\eta_{k+1} = \frac{1}{2}\eta_k$
9: Update $k = k + 1$
10: **end while**

To present the main convergence results of the Epoch-DSMD algorithm, we assume from now on that the constraint set has finite radius $R_\mathcal{W} = \max_{\mathbf{w},\mathbf{v}\in\mathcal{W}} \|\mathbf{w} - \mathbf{v}\|_2$, and denote the $\Phi(\cdot)$-diameter of $\mathcal{W}$ by

$$R_{\Phi,\mathcal{W}} := \left(\max_{\mathbf{w}\in\mathcal{W}} \Phi(\mathbf{w}) - \min_{\mathbf{w}\in\mathcal{W}} \Phi(\mathbf{w})\right)^{1/2}.$$

Theorem 2 *Under the conditions of Theorem 1, set the parameters in Algorithm 2 as* $\eta_1 = \frac{1}{\sigma_F}$ *and* $T_1 = 4$. *Then the final estimates* $\mathbf{w}_{i,1}^{k^\dagger+1}$ *enjoys a convergence rate of*

$$\sum_{i=1}^{m} \mathbb{E}[\|\mathbf{w}_{i,1}^{k^\dagger+1} - \mathbf{w}^*\|_2^2] \leq \frac{64\widehat{c}}{T}$$

where $k^\dagger = \lfloor \log_2\left(\frac{T}{4}+1\right)\rfloor$ *is the total number of epochs in Algorithm 2, and* $\widehat{c} = \max\left\{\frac{\sigma_F c_1 + 4c_2}{4\sigma_F^2\sigma_\Phi}, \frac{mR_{\Phi,\mathcal{W}}^2}{4\sigma_\Phi}\right\}$ *with* $c_1 = mG\left(\frac{2\alpha\beta}{1-\beta}+1\right)\left(mR_\mathcal{W} + \sum_{i=1}^m \mathbb{E}[\|\mathbf{w}_{i,1}^1\|_2]\right)$ *and* $c_2 = \frac{mG^2}{2} + \frac{2\alpha\beta m^2 G^2}{(1-\beta)\sigma_\Phi}$.

Before the proof of Theorem 2, we first provide the following lemma for a basic convergence result of Algorithm 2.

Lemma 4 *Under the conditions of Theorem 1. We have that, for any* $k \geq 1$,

$$\sigma_F \sum_{i=1}^{m} \mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^{k+1})] \leq \frac{c_1}{T_k} + c_2\eta_k + \frac{1}{\eta_k T_k}\sum_{i=1}^{m} \mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)].$$

Proof. See Appendix C. □

Armed with Lemma 4, we are in a position to present the proof of Theorem 2.

Proof. [Proof of Theorem 2] First, we show by induction that

$$\sum_{i=1}^{m} \mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)] \leq \max\left\{\frac{1}{\sigma_F}\left(\frac{c_1}{T_1} + c_2\eta_1\right), \frac{1}{4}mR_{\Phi,\mathcal{W}}^2\right\}\cdot 2^{-(k-3)}. \quad (12)$$

We now prove that it is true at $k = 1$. It follows from the definition of the Bregman divergence that

$$\sum_{i=1}^{m} D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^1) = \sum_{i=1}^{m}\left(\Phi(\mathbf{w}^*) - \Phi(\mathbf{w}_{i,1}^1) - \left\langle\nabla\Phi(\mathbf{w}_{i,1}^1), \mathbf{w}^* - \mathbf{w}_{i,1}^1\right\rangle\right) \quad (13)$$

utilizing the fact that $\mathbf{w}_{i,1}^k = \arg\min_{\mathbf{w}\in\mathcal{W}} \Phi(\mathbf{w})$ and applying the first order optimality condition for the term $\left\langle\nabla\Phi(\mathbf{w}_{i,1}^1), \mathbf{w}^* - \mathbf{w}_{i,1}^1\right\rangle$, we have

$$\left\langle\nabla\Phi(\mathbf{w}_{i,1}^1), \mathbf{w}^* - \mathbf{w}_{i,1}^1\right\rangle \geq 0$$

which, combined with (13), yields

$$\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^1)] \le \sum_{i=1}^{m}\mathbb{E}[\Phi(\mathbf{w}^*)-\Phi(\mathbf{w}_{i,1}^1)] \le mR_{\Phi,\mathcal{W}}^2. \quad (14)$$

Hence, the base of the induction holds. Assuming the bound (12) is true for k, we now show it holds for $k+1$ as well, by combining the fact that $T_k = T_1 2^{k-1}$ and $\eta_k = \eta_1 2^{-(k-1)}$ with (12):

$$\begin{aligned}
&\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^{k+1})] \\
&\le \frac{1}{\sigma_F}\left(\frac{c_1}{T_1}+c_2\eta_1\right)2^{-(k-1)} \\
&\quad + \frac{1}{\eta_1 T_1 \sigma_F}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^{k})] \\
&\le \max\left\{\frac{1}{\sigma_F}\left(\frac{c_1}{T_1}+c_2\eta_1\right), \frac{1}{4}mR_{\Phi,\mathcal{W}}^2\right\}\cdot 2^{-(k-1)} \\
&\quad + \frac{1}{4}\max\left\{\frac{1}{\sigma_F}\left(\frac{c_1}{T_1}+c_2\eta_1\right), \frac{1}{4}mR_{\Phi,\mathcal{W}}^2\right\}\cdot 2^{-(k-3)} \\
&= \max\left\{\frac{1}{\sigma_F}\left(\frac{c_1}{T_1}+c_2\eta_1\right), \frac{1}{4}mR_{\Phi,\mathcal{W}}^2\right\}\cdot 2^{-(k-2)} \quad (15)
\end{aligned}$$

where the first inequality holds by Lemma 4, and the second inequality follows from $\eta_1 = \frac{1}{\sigma_F}$ and $T_1 = 4$ and the induction hypothesis for $\sum_{i=1}^m \mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)]$. This shows that the relation (12) holds for all $k \ge 1$. Using the assumption on the strong convexity of function Φ, we have

$$\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)] \ge \sum_{i=1}^{m}\frac{\sigma_\Phi}{2}\mathbb{E}[\|\mathbf{w}^*-\mathbf{w}_{i,1}^k\|_2^2]. \quad (16)$$

This, along with (12), gives

$$\begin{aligned}
&\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}^*-\mathbf{w}_{i,1}^k\|_2^2] \\
&\le \frac{2}{\sigma_\Phi}\max\left\{\frac{1}{\sigma_F}\left(\frac{c_1}{T_1}+c_2\eta_1\right), \frac{1}{4}mR_{\Phi,\mathcal{W}}^2\right\}\cdot 2^{-(k-3)} \\
&= \max\left\{\frac{\sigma_F c_1 + 4c_2}{4\sigma_F^2\sigma_\Phi}, \frac{mR_{\Phi,\mathcal{W}}^2}{4\sigma_\Phi}\right\}\cdot 2^{-(k-4)}. \quad (17)
\end{aligned}$$

On the other hand, from the stopping criterion in Algorithm 2, the number of epochs is given by the largest value of k such that $\sum_{i=1}^k T_k \le T$, that is,

$$\sum_{i=1}^{k} T_1 2^{i-1} = T_1(2^k-1) \le T$$

which implies that the final epoch is given by

$$k^\dagger = \left\lfloor \log_2\left(\frac{T}{4}+1\right)\right\rfloor.$$

Applying the bound (17) to $k^\dagger+1$ we get

$$\begin{aligned}
\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}^*-\mathbf{w}_{i,1}^{k^\dagger+1}\|_2^2] &\le \widehat{c}\cdot 2^{-(k^\dagger+1-4)} \\
&\le \widehat{c}\cdot 2^{-\left(\log_2\left(\frac{T}{T_1}+1\right)-4\right)} \\
&\le \frac{16T_1\widehat{c}}{T} \\
&= \frac{64\widehat{c}}{T} \quad (18)
\end{aligned}$$

where in the second inequality we used the relation $k^\dagger + 1 \ge \log_2\left(\frac{T}{T_1}+1\right)$. The proof is complete. □

Remark 5 *Theorem 2 shows that Algorithm 2 converges at an $O(1/T)$ rate, matching that of the best previously known centralized stochastic subgradient algorithm (see, for example, Hazan & Kale, 2014; Rakhlin, Shamir, & Sridharan, 2012). In contrast to the epoch gradient descent algorithm in Hazan and Kale (2014), the Epoch-DSMD algorithm is in a distributed multi-agent setting, which makes the algorithm design and the associated convergence analysis results essentially different from those in Hazan and Kale (2014). To the best of our knowledge, our proposed algorithm is the first distributed algorithm that utilizes the idea of epoch gradient descent to develop a distributed stochastic mirror descent algorithm (different from the distributed stochastic subgradient algorithm in Tsianos and Rabbat (2012)). Moreover, it achieves the optimal rate of convergence for distributed stochastic strongly constrained convex optimization, without assuming smoothness of the objective functions, in contrast to the $O(\ln(T)/T)$ rate of convergence in Tsianos and Rabbat (2012) and Nedić and Olshevsky (2016) with Euclidean norms.*

Remark 6 *It is worth noting that the step size is keeping constant in each epoch, which means that each node need not coordinate the step size with those of its neighbors, but at the end of each epoch. Specifically, $O(\ln(T))$ coordinations of the step size among all the nodes are needed. This make the Epoch-DSMD algorithm easier to implement in a distributed setting, as opposed to the algorithms that use (coordinated) diminishing step size (see, e.g., Nedić & Olshevsky, 2016; Tsianos & Rabbat, 2012; Xi, Wu, & Khan, 2014).*

5 Simulation Results

In this section, we report some simulation results of the proposed algorithms. Specifically, we consider the following standard distributed estimation problem over sensor

networks (see, e.g., Rabbat & Nowak, 2004; Nedić & Olshevsky, 2016):

$$\min_{\mathbf{w}\in\mathcal{W}} \sum_{i=1}^{m} a_i \|\mathbf{w}-\mathbf{b}_i\|_2^2 \tag{19}$$

where $a_i \in \mathbf{R}_{++}$ and $\mathbf{b}_i \in \mathbf{R}^d$ are problem data known only to node i, and $\mathcal{W}$ is the constraint set known to all the nodes. We will consider two constraint sets, namely, (i) $\mathcal{W} = \{\mathbf{w} \in \mathbf{R}^d : \sum_{i=1}^d [\mathbf{w}]_i = 1, [\mathbf{w}]_i \geq 0, i \in [d]\}$; and (ii) $\mathcal{W} = \{\mathbf{w} \in \mathbf{R}^d : l_i \leq [\mathbf{w}]_i \leq u_i, i \in [d]\}$.

Implement the proposed algorithms over a random graph that consists of 50 nodes, which is generated as that of Xiao and Boyd (2004); the graph is shown in Fig. 1. In all cases we use the dimension of estimate $d = 10$, and the communication matrix generated according to the maximum-degree weights in Xiao and Boyd (2004). Note that the subgradient noises are random variables generated independent and identically distributed from the normal distribution $\mathcal{N}(0, \sigma I_{d\times d})$, where σ is the magnitude of the noises that will be specified in the sequel. The simulation results for the methods are based on the average of 50 realizations.

For the first constraint, i.e., the unit simplex, we use the distance-generating function $\Phi(\mathbf{w}) = \sum_{i=1}^d [\mathbf{w}]_i \ln[\mathbf{w}]_i$, and set the initial estimates as $\mathbf{w}_{i,1} = \mathbf{w}_{i,1}^1 = (\frac{1}{d}, \frac{1}{d}, \cdots, \frac{1}{d})^T$ for all i. In particular, the DSMD algorithm is just the distributed stochastic entropic descent algorithm (4)–(5). Note that the standard distributed stochastic projected subgradient algorithms in Nedić and Olshevsky (2016) and Tsianos and Rabbat (2012) involve a Euclidean projection onto the simplex at every step, which, as we have stated earlier, is equivalent to computing the solution of a d-dimensional nonlinear equation. This makes that our proposed algorithms are more favorable for this case. Figs. 2 and 3 provide respectively a plot of the average error (on the log-scale) versus the number of iterations T for three randomly selected nodes, for choices of $\sigma = 0.25$ and $\sigma = 0.5$. It can be seen that the proposed algorithms converge, and Epoch-DSMD converges faster than DSMD, which is in compliance with the convergence results obtained in the paper.

For the second constraint, i.e., a box constraint, we set the parameters in $\mathcal{W}$ as follows: $l_i = -1$, $u_i = 1$, for all i. We use the distance-generating function $\Phi(\mathbf{w}) = \frac{1}{2}\|\mathbf{w}\|_2$, and set the initial estimates as $\mathbf{w}_{i,1} = \mathbf{w}_{i,1}^1 = (0, 0, \cdots, 0)^T$ for all i. In this case, the DSMD algorithm becomes the standard distributed stochastic projected subgradient algorithm (2)–(3). Figs. 4 and 5 provide respectively a plot of the average error (on the log-scale) versus the number of iterations T for three randomly selected nodes, for choices of $\sigma = 0.25$ and $\sigma = 0.5$. It can be seen from Figs. 4 and 5 that the proposed Epoch-DSMD algorithm converges much faster than the standard distributed stochastic projected subgradient algorithm (see, e.g., Nedić & Olshevsky, 2016).

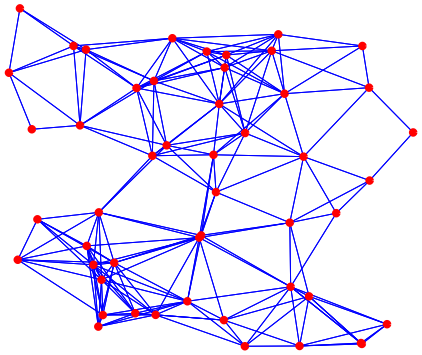

Fig. 1. The network.

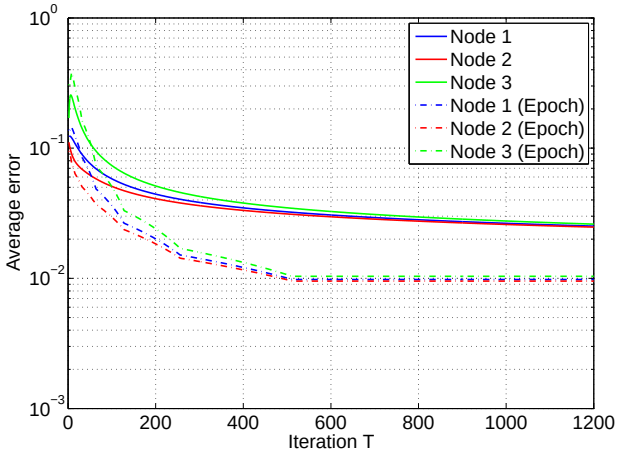

Fig. 2. The convergence of DSMD (in solid lines) and Epoch-DSMD (in dashed-dotted lines) for three randomly chosen nodes (the unit simplex constraint; noise magnitude $\sigma = 0.25$)

6 Conclusion

In this paper, we have studied the convergence rate problems for a distributed optimization problem over a switching multi-agent network. When the objective functions of agents are non-smooth and strongly convex (in the non-Euclidean sense), we have proposed two efficient non-Euclidean algorithms based on mirror descent. The first algorithm recovers the best previously known rate, and our second algorithm attains the optimal convergence rate. We have also reported some simulation results to illustrate the effectiveness of the proposed algorithms.

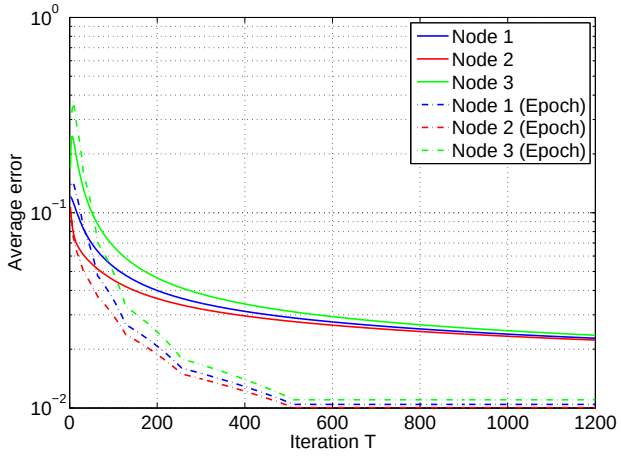

Fig. 3. The convergence of DSMD (in solid lines) and Epoch-DSMD (in dashed-dotted lines) for three randomly chosen nodes (the unit simplex constraint; noise magnitude $\sigma = 0.5$)

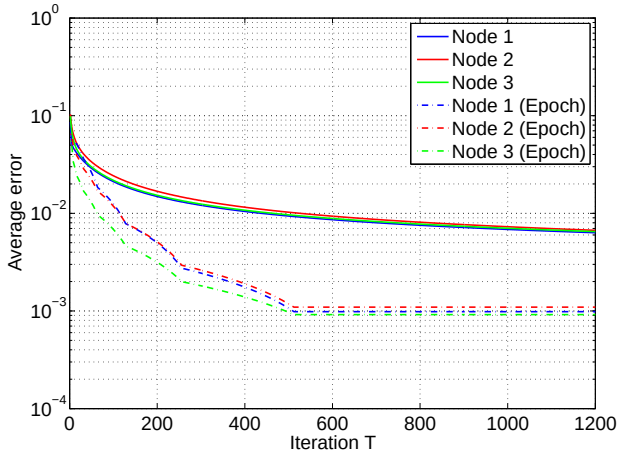

Fig. 4. The convergence of DSMD (in solid lines) and Epoch-DSMD (in dashed-dotted lines) for three randomly chosen nodes (the box constraint; noise magnitude $\sigma = 0.25$)

A Proof of Lemma 2

To simplify the notations, we denote

$$\overline{\mathbf{w}}_t = \frac{1}{m}\sum_{i=1}^{m} \mathbf{w}_{i,t} \tag{A.1}$$

$$\mathbf{r}_{i,t} = \mathbf{u}_{i,t+1} - \mathbf{w}_{i,t}. \tag{A.2}$$

We first provide an upper bound on the norm of the projection error $\mathbf{r}_{i,t}$ induced by the Bregman divergence. Using the first order optimality condition for the update formula for $\mathbf{u}_{i,t+1}$ and noting that $\nabla D_\Phi(\mathbf{u}_{i,t+1}||\mathbf{v}_{i,t}) = \nabla\Phi(\mathbf{u}_{i,t+1}) - \nabla\Phi(\mathbf{v}_{i,t})$, we obtain that, for all $\mathbf{v} \in \mathcal{W}$,

$$\langle \nabla\Phi(\mathbf{u}_{i,t+1}) - \nabla\Phi(\mathbf{v}_{i,t}), \mathbf{u}_{i,t+1} - \mathbf{v}\rangle \le 0.$$

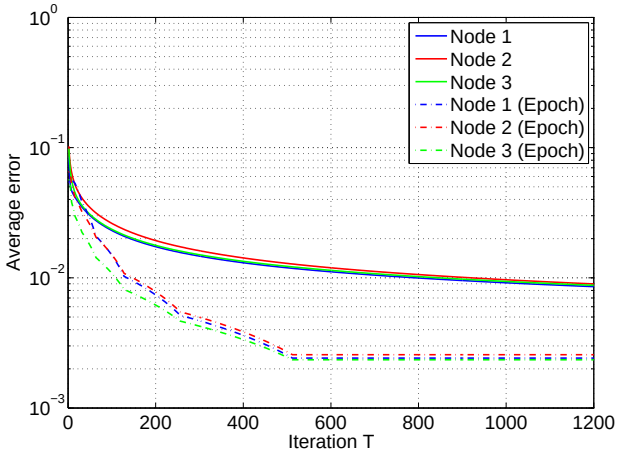

Fig. 5. The convergence of DSMD (in solid lines) and Epoch-DSMD (in dashed-dotted lines) for three randomly chosen nodes (the box constraint; noise magnitude $\sigma = 0.5$)

Since $\mathbf{w}_{i,t} \in \mathcal{W}$, by setting $\mathbf{v} = \mathbf{w}_{i,t}$ we have

$$\langle \nabla\Phi(\mathbf{u}_{i,t+1}) - \nabla\Phi(\mathbf{v}_{i,t}), \mathbf{u}_{i,t+1} - \mathbf{w}_{i,t}\rangle \le 0.$$

Substituting the update formula for $\nabla\Phi(\mathbf{v}_{i,t})$ into the preceding inequality yields

$$\begin{aligned}
&\eta_t \langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{u}_{i,t+1}\rangle \\
&\quad \ge \langle \nabla\Phi(\mathbf{u}_{i,t+1}) - \nabla\Phi(\mathbf{w}_{i,t}), \mathbf{u}_{i,t+1} - \mathbf{w}_{i,t}\rangle \\
&\quad \ge \sigma_\Phi \|\mathbf{u}_{i,t+1} - \mathbf{w}_{i,t}\|_2^2
\end{aligned}$$

where in the last inequality we used the fact that function Φ is σ_Φ-strongly convex, i.e., $\langle \nabla\Phi(\mathbf{w}) - \nabla\Phi(\mathbf{v}), \mathbf{w} - \mathbf{v}\rangle \ge \sigma_\Phi \|\mathbf{w} - \mathbf{v}\|_2^2$, for any two points $\mathbf{w}$ and $\mathbf{v}$. This leads to the following bound

$$\begin{aligned}
\mathbb{E}[\|\mathbf{r}_{i,t}\|_2] &= \mathbb{E}[\|\mathbf{u}_{i,t+1} - \mathbf{w}_{i,t}\|_2] \\
&\le \frac{\eta_t}{\sigma_\Phi}\mathbb{E}[\|\hat{\mathbf{g}}_{i,t}\|_2] \\
&\le \frac{G}{\sigma_\Phi}\eta_t
\end{aligned} \tag{A.3}$$

where in the last inequality we used the bound $\mathbb{E}[\|\hat{\mathbf{g}}_{i,t}\|_2] \le \left(\mathbb{E}[\|\hat{\mathbf{g}}_{i,t}\|_2^2]\right)^{1/2} \le G$, based on Assumption 3 and Jensen's inequality.

We now derive the general evolution of the average decision $\overline{\mathbf{w}}_t$:

$$\begin{aligned}
\overline{\mathbf{w}}_t &= \frac{1}{m}\sum_{i=1}^{m}\sum_{j=1}^{m}[\mathbf{P}(t-1)]_{ij}\mathbf{u}_{j,t} \\
&= \frac{1}{m}\sum_{j=1}^{m}\left(\sum_{i=1}^{m}[\mathbf{P}(t-1)]_{ij}\right)\mathbf{u}_{j,t}
\end{aligned}$$

$$= \frac{1}{m}\sum_{i=1}^{m} \mathbf{u}_{i,t}$$
$$= \frac{1}{m}\sum_{i=1}^{m} (\mathbf{w}_{i,t-1} + \mathbf{r}_{i,t-1})$$
$$= \overline{\mathbf{w}}_{t-1} + \frac{1}{m}\sum_{i=1}^{m} \mathbf{r}_{i,t-1}$$

where the third equality is based on the doubly stochasticity of the weight matrix $\mathbf{P}(t-1)$, and the second-to-last equality is based on the definition of $\mathbf{r}_{i,t}$. The above equality implies the following general iteration:

$$\overline{\mathbf{w}}_t = \overline{\mathbf{w}}_1 + \sum_{\ell=1}^{t-1} \frac{1}{m}\sum_{i=1}^{m} \mathbf{r}_{i,\ell}. \tag{A.4}$$

Similarly, we derive the recursive relation for $\mathbf{w}_{i,t}$, and obtain

$$\mathbf{w}_{i,t} = \sum_{j=1}^{m} [\mathbf{P}(t-1,1)]_{ij}\mathbf{w}_{j,1} + \sum_{\ell=1}^{t-1}\sum_{j=1}^{m} [\mathbf{P}(t-1,\ell)]_{ij}\mathbf{r}_{j,\ell}. \tag{A.5}$$

Combining the equations (A.4) and (A.5), we obtain, that for any $t \geq 2$,

$$\|\overline{\mathbf{w}}_t - \mathbf{w}_{i,t}\|_2 \leq \sum_{j=1}^{m} \left| [\mathbf{P}(t-1,1)]_{ij} - \frac{1}{m} \right| \cdot \|\mathbf{w}_{j,1}\|_2 + \sum_{\ell=1}^{t-1}\sum_{j=1}^{m} \left| [\mathbf{P}(t-1,\ell)]_{ij} - \frac{1}{m} \right| \cdot \|\mathbf{r}_{j,\ell}\|_2.$$

Taking the expectation and using Lemma 1 yields

$$\begin{aligned}
&\mathbb{E}[\|\overline{\mathbf{w}}_t - \mathbf{w}_{i,t}\|_2] \\
&\quad \leq \alpha\beta^{t-1}\delta_1(\mathbf{w}) + \sum_{\ell=1}^{t-1} \alpha\beta^{t-\ell} \sum_{j=1}^{m} \mathbb{E}[\|\mathbf{r}_{j,\ell}\|_2] \\
&\quad \leq \alpha\beta^{t-1}\delta_1(\mathbf{w}) + \sum_{\ell=1}^{t-1} \alpha\beta^{t-\ell} \sum_{j=1}^{m} \frac{G}{\sigma_\Phi}\eta_\ell \\
&\quad = \alpha\beta^{t-1}\delta_1(\mathbf{w}) + \frac{\alpha m G}{\sigma_\Phi} \sum_{\ell=1}^{t-1} \beta^{t-\ell}\eta_\ell
\end{aligned}$$

where $\delta_1(\mathbf{w}) = \sum_{i=1}^{m} \mathbb{E}[\|\mathbf{w}_{i,1}\|_2]$, and in the second inequality we used (A.3). Summing the preceding inequalities over $t = 1, \ldots, T$ gives

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[\|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2] \\
&\quad = \mathbb{E}[\|\mathbf{w}_{i,1} - \mathbf{w}_{j,1}\|_2] + \sum_{t=2}^{T} \mathbb{E}[\|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2] \\
&\quad \leq \delta_1(\mathbf{w}) + \sum_{t=2}^{T} 2\alpha\beta^{t-1}\delta_1(\mathbf{w}) + \frac{2\alpha m G}{\sigma_\Phi} \sum_{t=2}^{T}\sum_{\ell=1}^{t-1} \beta^{t-\ell}\eta_\ell \\
&\quad \leq \left(\frac{2\alpha\beta}{1-\beta} + 1\right)\delta_1(\mathbf{w}) + \frac{2\alpha\beta m G}{(1-\beta)\sigma_\Phi} \sum_{t=1}^{T} \eta_t
\end{aligned} \tag{A.6}$$

where the first inequality follows from $\mathbb{E}[\|\mathbf{w}_{i,1} - \mathbf{w}_{j,1}\|_2] \leq \mathbb{E}[\|\mathbf{w}_{i,1}\|_2] + \mathbb{E}[\|\mathbf{w}_{j,1}\|_2] \leq \delta_1(\mathbf{w})$ and $\mathbb{E}[\|\mathbf{w}_{i,t} - \mathbf{w}_{j,t}\|_2] \leq \mathbb{E}[\|\overline{\mathbf{w}}_t - \mathbf{w}_{i,t}\|_2] + \mathbb{E}[\|\overline{\mathbf{w}}_t - \mathbf{w}_{j,t}\|_2]$, and the last inequality is based on the following argument:

$$\sum_{t=2}^{T}\sum_{\ell=1}^{t-1} \beta^{t-\ell}\eta_\ell \leq \sum_{t=1}^{T-1} \eta_t \left(\sum_{\ell=1}^{T-1} \beta^\ell\right) \leq \frac{\beta}{1-\beta}\sum_{t=1}^{T-1} \eta_t \leq \frac{\beta}{1-\beta}\sum_{t=1}^{T} \eta_t.$$

Summing up the inequalities (A.6) over $i = 1, \ldots, m$ gives the final bound.

B Proof of Lemma 3

By referring to Step 3 in Algorithm 1, we can expand $\langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{w}^* \rangle$ as follows:

$$\begin{aligned}
&\langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{w}^* \rangle \\
&\quad = \frac{1}{\eta_t} \langle \nabla\Phi(\mathbf{v}_{i,t}) - \nabla\Phi(\mathbf{w}_{i,t}), \mathbf{w}^* - \mathbf{w}_{i,t} \rangle \\
&\quad = \frac{1}{\eta_t}(D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}) - D_\Phi(\mathbf{w}^*||\mathbf{v}_{i,t}) + D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t})) \\
&\quad \leq \frac{1}{\eta_t}(D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}) - D_\Phi(\mathbf{w}^*||\mathbf{u}_{i,t+1}) + D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t}))
\end{aligned} \tag{B.1}$$

where the second equality is based on the following property of the Bregman divergence:

$$\begin{aligned}
&\langle \nabla\Phi(\mathbf{u}) - \nabla\Phi(\mathbf{v}), \mathbf{w} - \mathbf{v} \rangle \\
&\quad = D_\Phi(\mathbf{w}||\mathbf{v}) - D_\Phi(\mathbf{w}||\mathbf{u}) + D_\Phi(\mathbf{v}||\mathbf{u})
\end{aligned} \tag{B.2}$$

and the last inequality follows from the Pythagorean Theorem for Bregman divergences, that is, $D_\Phi(\mathbf{w}^*||\mathbf{u}_{i,t+1}) \leq D_\Phi(\mathbf{w}^*||\mathbf{v}_{i,t})$, because $\mathbf{u}_{i,t+1}$ is the Bregman projection of $\mathbf{v}_{i,t}$ onto the convex set $\mathcal{W}$. Now, summing the inequalities (B.1) over all $i = 1, \ldots, m$, we obtain

$$\begin{aligned}
\sum_{i=1}^{m} \langle \hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t} - \mathbf{w}^* \rangle \leq \frac{1}{\eta_t}\sum_{i=1}^{m} &(D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}) \\
&- D_\Phi(\mathbf{w}^*||\mathbf{u}_{i,t+1}) + D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t})).
\end{aligned} \tag{B.3}$$

By Assumption 2(b) on the convexity of the Bregman divergence and Assumption 1(b) on the doubly stochas-

ticity of the weight matrix $\mathbf{P}(t)$, we obtain

$$\begin{aligned}\sum_{i=1}^{m} D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t+1}) &= \sum_{i=1}^{m} D_\Phi\left(\mathbf{w}^*\middle|\middle|\sum_{j=1}^{m}[\mathbf{P}(t)]_{ij}\mathbf{u}_{j,t+1}\right)\\ &\le \sum_{i=1}^{m}\sum_{j=1}^{m}[\mathbf{P}(t)]_{ij} D_\Phi\left(\mathbf{w}^*||\mathbf{u}_{j,t+1}\right)\\ &= \sum_{j=1}^{m} D_\Phi\left(\mathbf{w}^*||\mathbf{u}_{j,t+1}\right). \qquad \text{(B.4)}\end{aligned}$$

This, along with (B.3), yields

$$\begin{aligned}\sum_{i=1}^{m}\langle\hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t}-\mathbf{w}^*\rangle \le \frac{1}{\eta_t}\sum_{i=1}^{m}\big(&D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t})\\ &-D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t+1})+D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t})\big). \qquad \text{(B.5)}\end{aligned}$$

We are now in a position to bound $D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t})$. Owing to the definition of Bregman divergence we have

$$\begin{aligned}&D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t}) + D_\Phi(\mathbf{v}_{i,t}||\mathbf{w}_{i,t})\\ &\quad= \langle\nabla\Phi(\mathbf{w}_{i,t})-\nabla\Phi(\mathbf{v}_{i,t}), \mathbf{w}_{i,t}-\mathbf{v}_{i,t}\rangle\\ &\quad= \eta_t\langle\hat{\mathbf{g}}_{i,t}, \mathbf{w}_{i,t}-\mathbf{v}_{i,t}\rangle\\ &\quad\le \frac{1}{2}\|\eta_t\hat{\mathbf{g}}_{i,t}\|_2^2 + \frac{1}{2}\|\mathbf{w}_{i,t}-\mathbf{v}_{i,t}\|_2^2 \qquad \text{(B.6)}\end{aligned}$$

where the second equality follows from Step 3 in Algorithm 1, and the last inequality is based on the inequality that $\langle\mathbf{a},\mathbf{b}\rangle \le \frac{1}{2}\|\mathbf{a}\|_2^2 + \frac{1}{2}\|\mathbf{b}\|_2^2$, for any two vectors $\mathbf{a}$ and $\mathbf{b}$. In addition, using Assumption 2(a) on the Bregman divergence, we have

$$D_\Phi(\mathbf{v}_{i,t}||\mathbf{w}_{i,t}) \ge \frac{\sigma_\Phi}{2}\|\mathbf{v}_{i,t}-\mathbf{w}_{i,t}\|_2^2 \ge \frac{1}{2}\|\mathbf{v}_{i,t}-\mathbf{w}_{i,t}\|_2^2$$

which, combined with (B.6), yields

$$D_\Phi(\mathbf{w}_{i,t}||\mathbf{v}_{i,t}) \le \frac{\eta_t^2}{2}\|\hat{\mathbf{g}}_{i,t}\|_2^2. \qquad \text{(B.7)}$$

Combining the inequalities (B.5) and (B.7), we derive the conclusion.

C Proof of Lemma 4

By following the similar lines in Lemma 3, we have the following basic convergence result of Algorithm 2:

$$\begin{aligned}\sum_{i=1}^{m}\langle\hat{\mathbf{g}}_{i,t}^k, \mathbf{w}_{i,t}^k-\mathbf{w}^*\rangle \le \frac{\eta_k}{2}\sum_{i=1}^{m}\|\hat{\mathbf{g}}_{i,t}^k\|_2^2\\ +\frac{1}{\eta_k}\sum_{i=1}^{m}\left(D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k) - D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t+1}^k)\right). \text{(C.1)}\end{aligned}$$

Since η_k is keeping constant in each epoch k, by the fact that the term $\sum_{i=1}^m\left(D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k) - D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t+1}^k)\right)$ we can get a telescopic sum when summing over $t=1$ to $t=T_k$ and obtain

$$\begin{aligned}&\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\langle\hat{\mathbf{g}}_{i,t}^k, \mathbf{w}_{i,t}^k-\mathbf{w}^*\rangle]\\ &\quad= \sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\langle\mathbf{g}_i(\mathbf{w}_{i,t}^k), \mathbf{w}_{i,t}^k-\mathbf{w}^*\rangle]\\ &\quad\le \frac{\eta_k}{2}\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\|\hat{\mathbf{g}}_{i,t}^k\|_2^2] + \frac{1}{\eta_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)]\\ &\quad\le \frac{mG^2}{2}\eta_k T_k + \frac{1}{\eta_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)] \qquad \text{(C.2)}\end{aligned}$$

where the equality follows from the same reasoning as that of (6), and the last inequality from using the bound (11). Let us turn our attention to the left-hand side of the preceding inequality. By Assumption 2(b) on the strong convexity of function F_i, we have

$$\begin{aligned}&\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\langle\mathbf{g}_i(\mathbf{w}_{i,t}^k), \mathbf{w}_{i,t}^k-\mathbf{w}^*\rangle]\\ &\quad\ge \sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}\left[F_i(\mathbf{w}_{i,t}^k) - F_i(\mathbf{w}^*) + \sigma_F D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t})\right]\\ &\quad\ge \sum_{t=1}^{T_k}\mathbb{E}[F(\mathbf{w}_{j,t}^k)] - F(\mathbf{w}^*)\\ &\qquad+\sigma_F\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k)]\\ &\qquad-G\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,t}^k-\mathbf{w}_{j,t}^k\|_2] \qquad \text{(C.3)}\end{aligned}$$

where the last inequality is based on the same reasoning as that of (8). Using the fact that $\mathbf{w}^*$ is the minimizer of problem (1), we further have the following bound:

$$\begin{aligned}&\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\langle\mathbf{g}_i(\mathbf{w}_{i,t}^k), \mathbf{w}_{i,t}^k-\mathbf{w}^*\rangle]\\ &\quad\ge \sigma_F\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k)]\\ &\qquad-G\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,t}^k-\mathbf{w}_{j,t}^k\|_2]. \qquad \text{(C.4)}\end{aligned}$$

Combining the inequalities (C.2) and (C.4), and diving both sides by T_k, we get

$$\begin{aligned}
&\frac{\sigma_F}{T_k}\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k)]\\
&\quad\le \frac{mG^2}{2}\eta_k+\frac{1}{\eta_k T_k}\sum_{i=1}^{m}D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)\\
&\qquad+\frac{G}{T_k}\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,t}^k-\mathbf{w}_{j,t}^k\|_2]. \qquad \text{(C.5)}
\end{aligned}$$

We are left to bound the last term on the right-hand side of the preceding relation, i.e., the disagreement among all the nodes in the network. By replacing $\mathbf{P}(t)$ with $\mathbf{P}_k(t)$ in Algorithm 1 and following similar lines in Lemma 2, we immediate have that in epoch k, the disagreement among all the nodes satisfies

$$\begin{aligned}
&\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,t}^k-\mathbf{w}_{j,t}^k\|_2]\\
&\quad\le m\left(\frac{2\alpha\beta}{1-\beta}+1\right)\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,1}^k\|_2]+\frac{2\alpha\beta m^2G}{(1-\beta)\sigma_\Phi}\sum_{t=1}^{T_k}\eta_k\\
&\quad\le m\left(\frac{2\alpha\beta}{1-\beta}+1\right)\left(mR_{\mathcal{W}}+\widehat{\delta}_1(\mathbf{w})\right)+\frac{2\alpha\beta m^2G}{(1-\beta)\sigma_\Phi}\eta_k T_k \qquad \text{(C.6)}
\end{aligned}$$

where $\widehat{\delta}_1(\mathbf{w})=\sum_{i=1}^{m}\mathbb{E}[\|\mathbf{w}_{i,1}^1\|_2]$ and in the last inequality we used the compactness assumption of the set $\mathcal{W}$, that is, $\|\mathbf{w}_{i,1}^k\|_2=\|\mathbf{w}_{i,1}^k-\mathbf{w}_{i,1}^1+\mathbf{w}_{i,1}^1\|_2\le\|\mathbf{w}_{i,1}^k-\mathbf{w}_{i,1}^1\|_2+\|\mathbf{w}_{i,1}^1\|_2\le R_{\mathcal{W}}+\|\mathbf{w}_{i,1}^1\|_2$. Substituting the bound (C.6) into inequality (C.5) gives

$$\begin{aligned}
&\frac{\sigma_F}{T_k}\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k)]\\
&\quad\le mG\left(\frac{2\alpha\beta}{1-\beta}+1\right)\left(mR_{\mathcal{W}}+\widehat{\delta}_1(\mathbf{w})\right)\cdot\frac{1}{T_k}\\
&\qquad+\left(\frac{mG^2}{2}+\frac{2\alpha\beta m^2G^2}{(1-\beta)\sigma_\Phi}\right)\eta_k\\
&\qquad+\frac{1}{\eta_k T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)]\\
&\quad=\frac{c_1}{T_k}+c_2\eta_k+\frac{1}{\eta_k T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^k)]. \qquad \text{(C.7)}
\end{aligned}$$

Using the assumption on the convexity of the Bregman divergence D_Φ, the left-hand side can be further lower bounded by the following:

$$\begin{aligned}
&\frac{\sigma_F}{T_k}\sum_{t=1}^{T_k}\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k)]\\
&\quad=\sigma_F\sum_{i=1}^{m}\frac{1}{T_k}\sum_{t=1}^{T_k}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,t}^k)]\\
&\quad\ge\sigma_F\sum_{i=1}^{m}\mathbb{E}\left[D_\Phi\left(\mathbf{w}^*\middle\|\frac{1}{T_k}\sum_{t=1}^{T_k}\mathbf{w}_{i,t}^k\right)\right]\\
&\quad=\sigma_F\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\widehat{\mathbf{w}}_i^k)]\\
&\quad=\sigma_F\sum_{i=1}^{m}\mathbb{E}[D_\Phi(\mathbf{w}^*||\mathbf{w}_{i,1}^{k+1})] \qquad \text{(C.8)}
\end{aligned}$$

where the last equality follows from Step 7 in Algorithm 2. This, combined with (C.7), yields the conclusion.

References

Nedić, A., & Ozdaglar, A. (2009). Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, *54*(1), 48-61.

Kia, S. S., Cortés, J., & Martínez, S. (2015). Distributed convex optimization via continuous-time coordination algorithms with discrete-time communication. *Automatica*, *55*, 254–264.

Nedić, A., Ozdaglar, A., & Parrilo, P. A. (2010). Constrained consensus and optimization in multi-agent networks. *IEEE Transactions on Automatic Control*, *55*(4), 922–938.

Gharesifard, B., & Cortés, J. (2014). Distributed continuous-time convex optimization on weight-balanced digraphs. *IEEE Transactions on Automatic Control*, *59*(3), 781–786.

Ram, S. S., Nedić, A., & Veeravalli, V. V. (2010). Distributed stochastic subgradient projection algorithms for convex optimization. *Journal of Optimization Theory and Applications*, *147*(3), 516–545.

Lou, Y., Shi, G., Johansson, K. H., & Hong, Y. (2014). Approximate projected consensus for convex intersection computation: convergence analysis and critical error angle. *IEEE Transactions on Automatic Control*, *59*(7), 1722–1736.

Zhong, M., & Cassandras, C. G. (2010). Asynchronous distributed optimization with event-driven communication. *IEEE Transactions on Automatic Control*, *55*(12), 2735–2750.

Zhu, M., & Martínez, S. (2012). On distributed convex optimization under inequality and equality constraints. *IEEE Transactions on Automatic Control*, *57*(1), 151–164.

Lin, P., Ren, W., & Song, Y. (2016). Distributed multi-agent optimization subject to nonidentical constraints and communication delays. *Automatica*, *65*, 120–131.

Liu, S., Qiu, Z., & Xie, L. (2014). Continuous-time distributed convex optimization with set constraints. In *Proceedings of the 19th IFAC World Congress* (pp. 9762–9767).

Lu, J., Tang, C. Y., Regier, P. R., & Bow, T. D. (2011). Gossip algorithms for convex consensus optimization over networks. *IEEE Transactions on Signal Processing*, *56*(12), 2917–2923.

Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*, New York, NY: Cambridge University Press.

Rabbat, M., & Nowak, R. D. (2004). Distributed optimization in sensor networks. In *Proceedings of the International Conference on Information Processing in Sensor Networks* (pp. 20–27).

Yuan, D., Ho, D. W. C., & Xu, S. (2016). Regularized primal-dual subgradient method for distributed constrained optimization. *IEEE Transactions on Cybernetics*, *46*(9), 2109–2118.

Yuan, D., Ho, D. W. C., & Hong, Y. (2016). On Convergence rate of distributed stochastic gradient algorithm for convex optimization with inequality constraints. *SIAM Journal on Control and Optimization* (in press).

Liu Q., & Wang, J. (2015). A second-order multi-agent network for bound-constrained distributed optimization. *IEEE Transactions on Automatic Control*, *60*(12), 3310–3315.

Duchi, J., Agarwal, A., & Wainwright, M. (2012). Dual averaging for distributed optimization: convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, *57*(3), 592–606.

Jakovetić, D., Xavier, J., & Moura, J. M. F. (2014). Fast distributed gradient methods. *IEEE Transactions on Automatic Control*, *59*(5), 1131–1146.

Shi, W., Ling, Q., Wu, G., & Yin, W. (2015). EXTRA: an exact first-order algorithm for decentralized consensus optimization. *SIAM Journal on Optimization*, *25*(2), 944–966.

Chen, J., & Sayed, A. H. (2012). Diffusion adaptation strategies for distributed optimization and learning over networks. *IEEE Transactions on Signal Processing*, *60*(8), 4289–4305.

Chang, T. H., Nedić, A., & Scaglione, A. (2014). Distributed constrained optimization by consensus-based primal-dual perturbation method. *IEEE Transactions on Automatic Control*, *59*(6), 1524–1538.

Yi, P., Hong, Y., & Liu, F. (2015). Distributed gradient algorithm for constrained optimization with application to load sharing in power systems. *Systems & Control Letters*, *83*, 45–52.

Nedić, A., & Olshevsky, A. (2016). Stochastic gradient-push for strongly convex functions on time-varying directed graphs. *IEEE Transactions on Automatic Control* (in press, DOI: 10.1109/TAC.2016.2529285).

Beck, A., & Teboulle, M. (2003). Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, *31*, 167–175.

Nemirovski, A., Juditsky, A., Lan, G., & Shapiro, A. (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, *19*(4), 1574–1609.

Tsianos, K., & Rabbat, M. (2012). Distributed strongly convex optimization. In *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing* (pp. 593–600).

Xi, C., Wu, Q., & Khan, U. (2014). Distributed mirror descent over directed graphs. *Distributed mirror descent over directed graphs* (preprint, online at http://arxiv.org/abs/1412.5526).

Rabbat, M. (2015). Subgradient methods and consensus algorithms for solving convex optimization problems. In *Proceedings of the 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)* (pp. 517–520).

Hazan, E., & Kale, S. (2014). Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization. *Journal of Machine Learning Research*, *15*(1), 2489–2512.

Rakhlin, A., Shamir, O., & Sridharan, K. (2012). Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the 29th International Conference on Machine Learning* (pp. 449–456).

Nedić, A., Olshevsky, A., Ozdaglar, A., & Tsitsiklis, J. N. (2008). Distributed subgradient methods and quantization effects. In *Proceedings of the IEEE conference on decision and control* (pp. 4177–4184).

Li, J., Chen, G., Dong, Z., & Wu, Z. (2016). Distributed mirror descent method for multi-agent optimization with delay. *Neurocomputing*, *177*, 643–650.

Xiao, L., & Boyd, S. (2004). Fast linear iterations for distributed averaging. *Systems & Control Letters*, *53*(1), 65–78.